Journal of Power Sources 239 (2013) 705–710

Contents lists available at SciVerse ScienceDirect

Journal of Power Sources

journal homepage: www.elsevier.com/locate/jpowsour

Particle filter for state of charge and state of health estimation for lithium–iron phosphate batteries☆

Simon Schwunk*, Nils Armbruster, Sebastian Straub, Johannes Kehl, Matthias Vetter

Fraunhofer Institut fuer Solare Energiesysteme, Heidenhofstrasse 2, D-79110 Freiburg, Germany

HIGHLIGHTS

► First estimation of state of charge and health for a LFP battery with a particle filter.
► Particle filter allowing any probability distribution for state of charge and health.
► Modelling of the open-circuit voltage hysteresis by multimodal probability functions.
► Validated for applications like electric vehicles and photovoltaic off-grid power supply.
► State estimation validated for new and aged batteries.

ARTICLE INFO

Article history:
Received 6 September 2012
Received in revised form
12 October 2012
Accepted 15 October 2012
Available online 26 October 2012

Keywords:
Lithium-ion
Stochastic filter
State-of-charge
State-of-health
Monte Carlo
Particle filter

ABSTRACT

The paper presents a new approach for state estimation of lithium–iron phosphate batteries. Lithium–iron phosphate/graphite batteries are very intricate in state of charge estimation since the open circuit voltage characteristic is flat and ambiguous. The characteristic is ambiguous because open circuit voltages are different if one charges or discharges the battery. These properties also hinder state of health estimation. Therefore conventional approaches like Kalman filtering which represents a state by only the mean and the variance of a Gaussian probability density function tend to fail. The particle filter presented here overcomes the problem by using Monte Carlo sampling methods which are able to represent any probability density function. The ambiguities can be modelled stochastically and complex models dealing with hysteresis can be avoided. The state of health estimation employs the same framework and takes the estimated state of charge as input for estimating the battery's state of health. The filter was developed for A123 lithium–iron phosphate batteries. For validation purposes user profiles for batteries in different ageing states like electric vehicles and off-grid power supply applications were generated at a battery testing system. The results show good accuracy.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

Lithium-ion batteries are the most favoured battery technology in many upcoming applications today, may it be for electric vehicles or for storing renewable energy. Classical lithium-ion technology though suffers from safety problems and cost of raw materials as e.g. cobalt [1]. Therefore researchers are working on alternative materials like phosphate based batteries [2]. This is a very interesting class of cathode materials especially for safety reasons. But phosphate materials tend to have very flat open circuit voltages with a hysteresis between charging and discharging. For $LiFePO_4$, the most widely spread of these materials, this characteristic is shown in Fig. 1. During partial cycling even micro hysteresis appear like one can also see for nickel based batteries [3,4]. This is a problem for system engineering which needs to estimate the state of charge based on the terminal voltage. This paper will introduce a framework called particle filter which overcomes this problem by modelling this behaviour stochastically.

First the method will be introduced, afterwards the method will be proven based on real measurement data from the batteries. In the conclusions the results are assessed and further works will be described. For this paper 2.3 Ah batteries from A123 are used having the type designation ANR26650M1.

2. Description of the parallel particle filter

The most popular filter within the family of Bayesian filters is the Kalman filter [5–12]. The Kalman filter is an analytical solution

☆ Paper was presented at the 13th Ulm Electrochemical Talks.
* Corresponding author. Tel.:+49 761 4588 5219; fax: +49 761 4588 9000.
E-mail address: simon.schwunk@ise.fraunhofer.de (S. Schwunk).

0378-7753/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jpowsour.2012.10.058

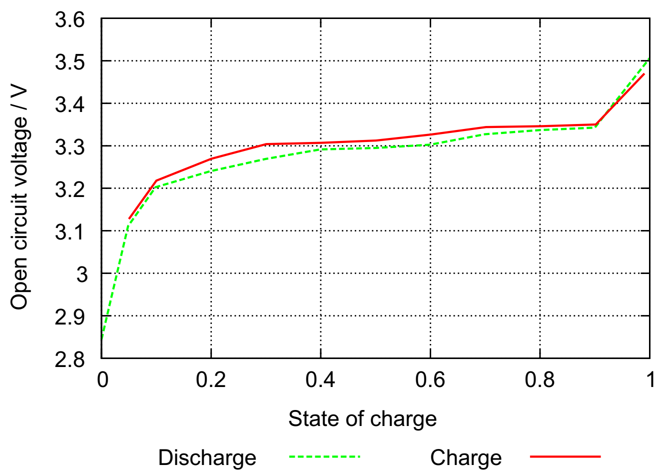

Fig. 1. Open circuit voltage curves of a lithium–iron phosphate battery during charging and discharging of the battery. One can clearly see the very flat characteristic and the hysteresis between the two curves.

of the Bayesian filter for Gaussian distributions. By employing Monte Carlo sampling techniques the particle filter offers the possibility to deal with any kind of distribution by approximating the respective probability density function by a set of particles or samples. For $LiFePO_4$-batteries with their hysteresis in the open circuit voltage this offers the possibility to use multimodal distributions modelling this behaviour stochastically.

2.1. Recursive Bayesian filtering

A recursive Bayesian filter estimates the state of dynamic systems online. It takes into account the probability for reaching a state x_t with respect to the states estimated in time steps $x_{1:t-1}$ before and the probability for observing a certain quantity z_t when being in a certain state x_t. The system is described by three quantities:

1. State x_t: This is the state of the system at time t, which shall be estimated. The state can neither be observed nor measured.
2. Input u_t: The input will change the system's state over time and can be measured.
3. Output z_t: The output is in some correlation with the system's state x_t enabling at least a rough estimation of the state. The quantity can be observed and measured.

The Bayesian filter estimates the following probability density function based on the quantities explained before:

$$P(x_t|z_{1:t}, u_{1:t-1}) = P(x_t|z_t, z_{1:t-1}, u_{1:t-1}) \tag{1}$$

As one applies the Bayes' theorem [13] one gets the following equation. The denominator normalising the function to one is written as η.

$$P(x_t) = \frac{P(z_t|x_t, z_{1:t-1}, u_{1:t-1}) \cdot P(x_t|z_{1:t-1}, u_{1:t-1})}{P(z_t|z_{1:t-1}, u_{1:t-1})} \tag{2}$$

$$= \eta^{-1} P(z_t|x_t, z_{1:t-1}, u_{1:t-1}) \cdot P(x_t|z_{1:t-1}, u_{1:t-1}) \tag{3}$$

Now the equation is marginalised over x_{t-1} and one assumes a Markov chain [13]. Therefore only the quantities u_{t-1} and z_t and not the former values have an effect on the calculated probability of the current state:

$$P(x_t) = \eta^{-1} P(z_t|x_t) \cdot \int P(x_t|x_{t-1}, u_{t-1}) \cdot P(x_{t-1}) \mathrm{d}x_{t-1} \tag{4}$$

The probability density function $P(x_{t-1})$ describes the state's probability density function in the time step before. The function $P(x_t|x_{t-1}, u_{t-1})$ takes the influence of the inputs u_{t-1} for the progression of the system from the state x_{t-1} to the state x_t into account. The function $P(z_t|x_t)$ represents the probability for observing the measurement z_t given a certain state x_t.

2.2. The particle filter

Since for most probability density functions there is no easy solution of this equation, techniques were developed for coping with more complex distributions than the Gaussian distributions of the Kalman filter. In Refs. [13,14] methods are introduced, which employ Monte Carlo methods for representing the distributions. The state distribution is represented by a set of samples, therefore being able to represent any distribution though of course with accuracy limited by the number of particles employed. High probabilities are represented by a huge number of particles in a certain area, low probabilities by a low number or no particles in an area.

The algorithm is performed in three steps. The first step is the state transition. In this step the influence of the input u_{t-1} on each sample is calculated. The uncertainty of this step and the measurement errors of the input are taken into account by the addition of some noise on each sample during state transition. The result is the probability density function $P(x_t|x_{t-1}, u_{t-1})$. Over time there will be a diffusion of the samples increasing the variance of the estimation.

The second and the third step assess the plausibility of a state x_t when observing a measurement z_t. This leads to decreasing the variance and limiting the diffusion of the particles.

In the second step the samples are weighted. According to the observed measurement value z_t a weight w_t^k of each sample s_t^k is calculated via a probability density function $P(z_t|x_t)$. Afterwards the overall weight of all samples is normalised to one.

The third step is resampling all the samples according to their weights. After this step all samples have the same weight again. In Ref. [13] several sampling methods are introduced that might be employed. For low calculation effort low variance resampling is employed [15].

2.3. Application of the particle filter

In this paper the particle filter is used for estimating the two most important quantities: the state of charge of the battery and the state of health of the battery. The state of charge is defined as the amount of charge in the battery divided by its capacity:

$$\mathrm{SOC} = \frac{Q}{C_{\mathrm{batt}}} \tag{5}$$

The state of health of the battery is a normalised value defined as the battery's actual capacity C_{batt} divided by its nominal capacity C_r:

$$\mathrm{SOH} = \frac{C_{\mathrm{batt}}}{C_r} \tag{6}$$

Each of the quantities is represented by a set of particles of size N_{samples}. The sizes of the sets may differ. For both quantities a separate particle filter is implemented and they exchange the averages of the samples representing the quantity. The state of charge filter takes the estimated state of health as a parameter to its

model during state transition, which is only updated in long time intervals. The state of health filter takes the state of charge as the output z_t of the system, which can be observed for weighting the samples, which represent the state of health.

2.4. Estimating the state of charge

The state of charge is the state x_t of the system. Before the particle filter starts working all samples need to be initialized. Because the actual state of charge is generally unknown, all samples are equally distributed over the complete state of charge range. During state transition the following equation calculates the influence of the battery current I_{batt} under consideration of a noise ε representing the modelling and measurement uncertainties:

$$s^k_{\mathrm{SOC},t} = s^k_{\mathrm{SOC},t-1} + \frac{\left(I_{\mathrm{batt}} + \varepsilon^k\right)\Delta t}{\mathrm{SOH} \cdot C_r} \tag{7}$$

The noise ε^k is sampled from a suitable distribution. For avoiding too little diffusion a Cauchy Lorentz distribution, which has very fat tails compared to a Gaussian distribution, is employed:

$$p(x) = \frac{1}{\pi} \cdot \frac{s}{s^2 + (x - x_0)^2} \tag{8}$$

For sampling from this distribution one draws samples c uniformly [0,1] and solves the following integral for receiving the nonuniformly distributed samples y [13]:

$$c = \int_{-\mathrm{inf}}^{y} p(x)\mathrm{d}x \tag{9}$$

For a Cauchy Lorentz distribution the following equation transforms the uniformly distributed into Cauchy distributed samples:

$$y = x_0 + s \cdot \tan(\pi \cdot (c - 0.5)) \tag{10}$$

The addition of noise leads to a diffusion of the samples making the estimation more inaccurate over time. For decreasing the diffusion again a model is employed making it possible to compare the samples to a measured value and checking each of them for plausibility. In case of the state of charge estimation a quasistationary battery model is employed. Since this is a $LiFePO_4$ battery the hysteresis makes the possible state of charge values ambiguous, resulting in several possibilities for the correct state of charge value.

This is stochastically modelled by the function $P(z_t|x_t)$, in this case a bimodal Cauchy distribution. Two possibilities for the expected terminal voltage for each sample, $V_{\mathrm{discharge}}$ and V_{charge} are calculated by the following equations for the discharge and the charge case of the open circuit voltage:

$$V^k_{\mathrm{discharge},t} = \mathrm{OCV}_{\mathrm{discharge}}\left(s^k_{\mathrm{SOC},t}\right) + R_i\left(s^k_{\mathrm{SOC},t}, T, I_{\mathrm{batt}}\right) \cdot I_{\mathrm{batt}} \tag{11}$$

$$V^k_{\mathrm{charge},t} = \mathrm{OCV}_{\mathrm{charge}}\left(s^k_{\mathrm{SOC},t}\right) + R_i\left(s^k_{\mathrm{SOC},t}, T, I_{\mathrm{batt}}\right) \cdot I_{\mathrm{batt}} \tag{12}$$

Both voltages are compared to the measured terminal voltage and two probabilities are determined by the Cauchy distribution and added to form the weight of each sample:

$$w^k_{\mathrm{SOC},t} = \frac{1}{\pi} \cdot \frac{s}{s^2 + \left(V^k_{\mathrm{discharge},t} - V_{\mathrm{meas}}\right)^2} + \frac{1}{\pi} \cdot \frac{s}{s^2 + \left(V^k_{\mathrm{charge},t} - V_{\mathrm{meas}}\right)^2} \tag{13}$$

The long tails avoid sample impoverishment and lead to higher stability and robustness of the algorithm. This results in a set of couples consisting of a sample and its probability. These probabilities are normalised to one and then resampled using the low variance resampling.

The low variance resampling puts all the samples in one line assigning each sample a length according to its weight. For starting the resampling process a random starting value is chosen within the interval $[0,1/N)$ in which N is the number of samples. Afterwards the algorithm goes forward with step size $1/N$ in $N - 1$ steps and chooses at each step the corresponding sample. Thereby exactly the same amount of new samples as the old sample set is generated and samples with small weight are neglected and samples with a high weight are chosen several times. The algorithm is illustrated in Fig. 2.

Since a bimodal Cauchy distribution is used also the Monte Carlo representation can be multimodal as is illustrated in Fig. 3. For the state estimation the mean of the particles is calculated and used as estimation.

2.5. Implementation of the state of health filter

The state of health is the state of the system x_t. The state of health filter assumes no inputs that influence the capacity of the battery, resulting in $P(x_t|x_{t-1},u_{t-1}) = P(x_t|x_{t-1})$. A very slow diffusion of the particles is assumed for modelling the capacity change over time. For the diffusion a noise drawn from a Cauchy distribution is added to each particle:

$$s^k_{\mathrm{SOH},t} = s^k_{\mathrm{SOH},t-1} + \varepsilon^k \tag{14}$$

For reacting on the diffusion processes the output of the state of charge filter is used to calculate the measurement value $\Delta\mathrm{SOC}_{\mathrm{measured}}$ and compares it to the integrated charge flown into and out of the battery Q_{step}. Therefore the difference of state of charge between two points in time is the measurement value: $z_t = \Delta\mathrm{SOC}_t$. Each of the particle is weighted by a Cauchy distribution

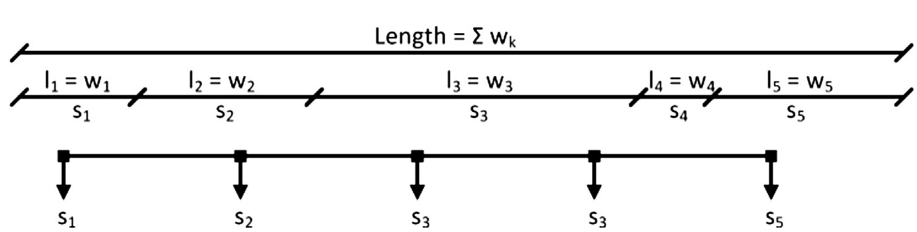

Fig. 2. Low variance resampling for generating a new unweighted sample set from a weighted sample set.

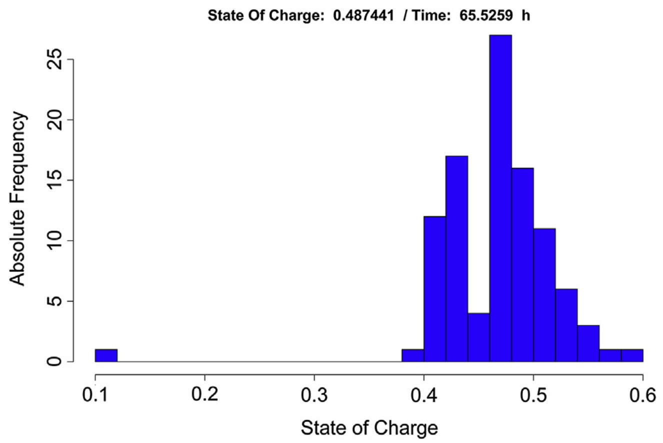

Fig. 3. Histogram of the particles during cycling in a medium and ambiguous state of charge. One can clearly see, how a bimodal distribution evolves due to the modelled hysteresis of the open circuit voltage.

and assigned a weight accordingly. The expected state of charge change is calculated for each particle describing the state:

$$\Delta \mathrm{SOC}_t^k = \frac{Q_{\mathrm{step}}}{s_{\mathrm{SOH},t}^k \cdot C_r} \tag{15}$$

The weight for each of these particles is then assigned by a Cauchy distribution:

$$w_{\mathrm{SOH}}^k = \frac{1}{\pi} \cdot \frac{s}{s^2 + \left(\Delta \mathrm{SOC}_t^k - \Delta \mathrm{SOC}_{\mathrm{measured}}\right)^2} \tag{16}$$

After assigning a weight to each particle the resampling step is performed with the low variance resampling method and a new unweighted sample set is generated. For estimating the state of health the mean of the particles is generated.

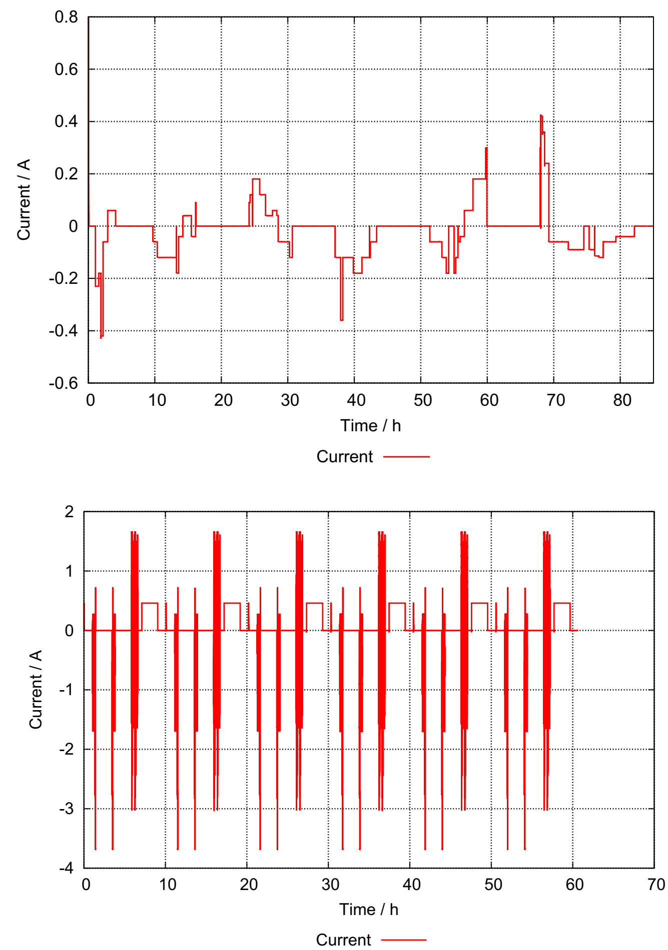

Fig. 4. Current profiles for a photovoltaic off-grid power supply and an electric vehicle application as applied here for validating the state of charge and state of health estimations.

3. Validation

3.1. Measurements

For validating the algorithm current profiles were run on a battery cycler which shall represent specific applications. Two applications were chosen. First a profile for a photovoltaic off grid power supply is generated with low currents but only little breaks in the consumption. Second a profile is generated with more breaks but higher charge and discharge currents, which are based on standard driving cycles, representing an electric vehicle application. The photovoltaic profile was performed on a new battery having the state of health 1.0 and the electric vehicle profile on an already aged battery having the state of health 0.939. The ageing states were determined in capacity tests before profiles were run. The current profiles can be seen in Fig. 4.

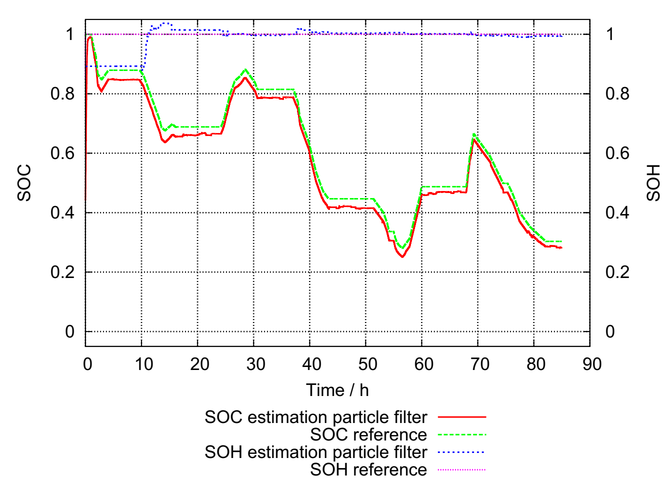

Fig. 5. Results for a current profile in a photovoltaic off-grid power supply. The filter finds both state of charge and state of health correctly. The battery has got a state of health of 100%. The reference for the state of charge was generated by fully charging the battery for knowing the exact state of charge at the beginning of the test and measuring and integrating the current afterwards. For the state of health a capacity test was performed before applying the off grid power supply profile on the battery. The references are therefore both measured data.

3.2. Photovoltaic profile

In Fig. 5 you see the reference state of charge and state of health for a lithium–iron phosphate battery with state of health 100%. At the beginning the sample set for the state of charge estimation is equally distributed in the interval [0,1] resulting in starting value of about 0.5 for the state of charge estimation. The sample set for the state of health estimation is equally distributed in the interval [0.35,1.35] resulting in a starting value of about 0.85 for the state of

health estimation. The same value is used as starting value for the state of health used in the state of charge estimation.

As can be seen the state of health estimation only starts after the state of charge has found a reasonable value. For achieving that a starting time for the state of health filter is introduced. The state of charge filter quickly finds the correct value and is able to track it over the whole time, while still doing some corrections as can be seen during the rest phases. This guarantees, that the state of charge estimation is no pure Ah counter.

The filter tracks both estimated states very well over time, is stable and there is hardly no noise on the quantities. This proves that the stochastic modelling of the hysteresis is sufficient to cope with the ambiguities of the hysteresis and the resulting uncertainties for the state of charge estimation.

For proving that the algorithm is not in need of starting with a full charge, in another run the algorithm is started later at a lower state of charge. As can be seen clearly in Fig. 6, the state of charge is found quickly and tracked well over time.

3.3. Electric vehicle profile

The filter is parameterised exactly the same way for the electric vehicle profile as for the photovoltaic profile. An aged battery with state of health 93.9% is used and for this aged battery results are also very good and robust, though the underlying battery model is only quasistationary and developed for a new battery. Therefore the ageing of the battery resulting in an increase in inner resistance is not captured by the model.

The results for state of charge and state of health from the validation are shown in Fig. 7. The state of charge is tracked very well over the whole time providing a good basis for the state of health estimation. Therefore also the correct state of health is found quickly and then stably tracked over time.

3.4. Non-deterministic behaviour

In Fig. 8 the non-deterministic behaviour of the particle filter is shown. As can be seen each state of charge estimation and state of health estimation is different, though the data, the algorithm and the parameters are the same. The only difference is the seed for the function generating the equally distributed random number. Those are used for the noise added during the calculation the effect of the

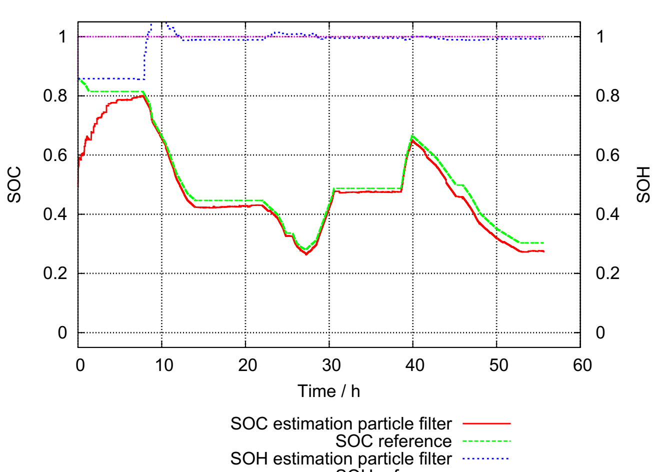

Fig. 6. The particle filter is started later on the same data set as for Fig. 5, starting therefore only at a state of charge of 80%. As can be seen clearly no full charge is needed for a good state estimation.

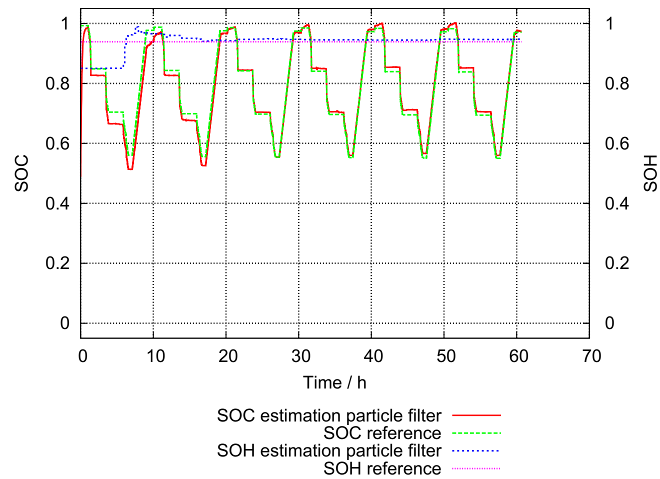

Fig. 7. Results for applying the particle filter to an electric vehicle profile. The battery has only got a state of health of 93.9%. The reference for the state of charge was generated by completely charging the battery to know the exact state of charge at the beginning of the test and measuring and integrating the current afterwards. For the state of health a capacity test was performed before applying the electric vehicle profile on the battery. Both references are therefore measured data.

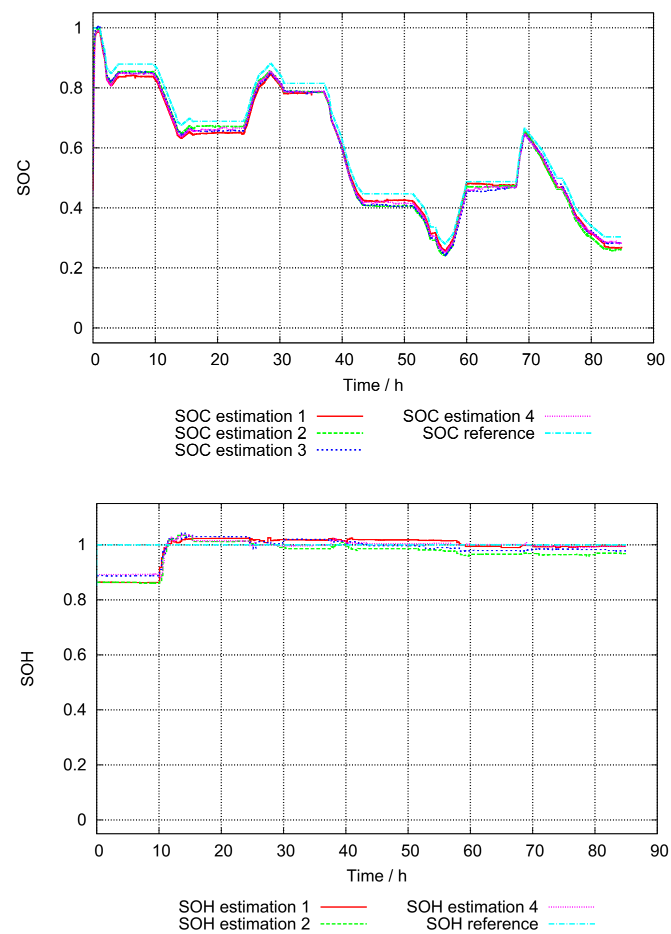

Fig. 8. Different runs of the identical particle filter for the same data, in this case for an off-grid power supply. As can be seen clearly the result differs from run to run due to the random numbers drawn by the Monte Carlo method.

inputs on the states and for choosing the starting value for the low variance resampling.

4. Conclusions

A framework for dealing with difficult and ambiguous batteries like $LiFePO_4$/graphite batteries was presented. The ambiguous range of the open circuit voltage is stochastically modelled opening up possibilities of multimodal estimations for the battery's state of charge by a Monte Carlo representation of its state distribution with samples. Thereby the algorithm is both able to find a correct value at the start and track the state of charge correctly. This forms the basis for the good state of health estimation.

The main advantage of this stochastic modelling compared to a physical approach for tracking the hysteresis behaviour is a quick modelling during development time, no starting state for the hysteresis needed and a possibility for indicating several plausible states with their probabilities. The last issue is especially interesting during the start up phases.

The algorithm reaches a high accuracy in the examples shown in the paper. The error values for state of charge and state of health estimation settle at very low levels.

Future work will focus on analysing and increasing robustness of the method, on extracting more information out of the sample sets like error boundaries and estimation of correctness and analysing the long term behaviour of the algorithm during ageing.

For increasing accuracy and accelerating the determination of the correct state of charge one could use more complex impedance models instead of the more simplified battery model used in this paper. In general this algorithm could be a possible replacement for the different kinds of Kalman filters used today as state of the art for state estimation.

References

[1] J. Goodenough, Journal of Power Sources 174 (2) (2007) 996–1000.
[2] W.F. Howard, R. Spotnitz, Journal of Power Sources 165 (2007) 887–891.
[3] Marc Thele, Oliver Bohlen, Dirk Uwe Sauer, Journal of Power Sources 175 (2008) 635–643.
[4] J. Gerschler, D.U. Sauer, in Proceeding of the 24th Electric Vehicle Symposium, Stavanger, Norway, 2009.
[5] R.E. Kalman, Journal of Basic Engineering 85 (1960) 35–45.
[6] S. Lee, J. Kim, J. Lee, B. Cho, Journal of Power Sources 185 (2) (2008) 1367–1373.
[7] M. Urbain, S. Rael, B. Davat, P. Desprez, in: IEEE-PESC'07, 2007, pp. 2804–2810.
[8] S. Piller, M. Perrin, A. Jossen, Journal of Power Sources 96 (1) (2001) 113–120.
[9] J. Lee, O. Nam, B. Cho, Journal of Power Sources 174 (7) (2007) 9–15.
[10] G.L. Plett, Journal of Power Sources 134 (2) (2004) 252–261.
[11] G.L. Plett, Journal of Power Sources 134 (2) (2004) 262–276.
[12] G.L. Plett, Journal of Power Sources 134 (2) (2004) 277–292.
[13] C.M. Bishop, Pattern Recognition and Machine Learning, Springer, 2006.
[14] Michael Isard, Andrew Blake, International Journal of Computer Vision 29 (1) (1998) 5–28.
[15] J. Carpenter, P. Clifford, P. Fearnhead, IEE Proc. Radar Sonar Navig. 146 (1) (1999) 2–7.